

**REPLICATING KERNEL DATA THAT IS READ MOSTLY AND WRITE
RARELY**

ABSTRACT

A distributed computer system and a method for replicating data in a distributed computer system. The system includes a plurality of processing nodes. A processing node comprises at least one processor and at least one local memory, where the local memory is in communication with each processing node. The system also includes maintenance software to determine whether data is read substantially more frequently than it is written and to replicate the data that is read substantially more frequently than it is written among the plurality of processing nodes. The method includes reviewing classes of data, identifying whether at least a portion of data of a certain class used by the processing nodes is read substantially more frequently than it is written and replicating copies of the data of that class in the local memories.